arXiv:1603.03678v2 [stat.ML] 22 May 2016

Nonstationary Distance Metric Learning

Kristjan Greenewald
EECS Department
University of Michigan
Ann Arbor, MI 48109
greenewk@umich.edu

Stephen Kelley
MIT Lincoln Laboratory
Lexington, MA 02420
stephen.kelley@ll.mit.edu

Alfred O. Hero III
EECS Department
University of Michigan
Ann Arbor, MI 48109
hero@umich.edu

Abstract

Recent work in distance metric learning has focused on learning transformations of data that best align with provided sets of pairwise similarity and dissimilarity constraints. The learned transformations lead to improved retrieval, classification, and clustering algorithms due to the better adapted distance or similarity measures. Here, we introduce the problem of learning these transformations when the underlying constraint generation process is nonstationary. This nonstationarity can be due to changes in either the ground-truth clustering used to generate constraints or changes to the feature subspaces in which the class structure is apparent. We propose and evaluate COMID-SADL, an adaptive, online approach for learning and tracking optimal metrics as they change over time that is highly robust to a variety of nonstationary behaviors in the changing metric. We demonstrate COMID-SADL on both real and synthetic data sets and show significant performance improvements relative to previously proposed batch and online distance metric learning algorithms.

1 Introduction

The effectiveness of many machine learning and data mining algorithms depends on an appropriate measure of pairwise distance between data points that accurately reflects the learning task, e.g., prediction, clustering or classification. The kNN classifier, K-means clustering, and the Laplacian-SVM semi-supervised classifier are examples of such *distance-based* machine learning algorithms. In settings where there is clean, appropriately-scaled spherical Gaussian data, standard Euclidean distance can be utilized. However, when the data is heavy tailed, multimodal, or contaminated by outliers, observation noise, or irrelevant or replicated features, use of Euclidean inter-point distance can be problematic, leading to bias or loss of discriminative power.

To reduce bias and loss of discriminative power of distance-based machine learning algorithms, data-driven approaches for optimizing the distance metric have been proposed. These methodologies, generally taking the form of dimensionality reduction or data "whitening", aim to utilize the data itself to learn a transformation of the data that embeds it into a space where Euclidean distance is appropriate. Examples of such techniques include Principal Component Analysis [3], Multidimensional Scaling [12], covariance estimation [12, 3], and manifold learning [17]. Such unsupervised methods do not exploit human input on the distance metric, and they overly rely on prior assumptions, e.g., local linearity or smoothness.

In distance metric learning one seeks to learn transformations of the data that are well matched to a particular task specified by the user. Point labels or constraints indicating point similarity or dissimilarity are used to learn a transformation of the data such that similar points are "close" to one another and dissimilar points are distant in the transformed space. Learning distance metrics in this manner allows a more precise notion of distance or similarity to be defined that is related to the task at hand.

29th Conference on Neural Information Processing Systems (NIPS 2016), Barcelona, Spain.

Many supervised and semi-supervised distance metric learning approaches have been developed [15]. This includes online algorithms [16] with regret guarantees for situations where similarity constraints are received sequentially.

This paper proposes a new method that provides distance metric tracking. Specifically, we suppose the underlying ground-truth (or optimal) distance metric from which constraints are generated is evolving over time, in an unknown and potentially nonstationary way. We propose an adaptive, online approach to track the underlying metric as the constraints are received. Our algorithm, which we call COMID-Strongly Adaptive Dynamic Learning (COMID-SADL) is inspired by recent advances in composite objective mirror descent for metric learning [9] (COMID) and the Strongly Adaptive Online Learning (SAOL) framework proposed in [6]. We prove strong bounds on the dynamic regret of every subinterval, guaranteeing strong adaptivity and robustness to nonstationary metric drift such as discrete shifts, slow drift with a nonstationary drift rate, and combinations thereof.

1.1 Related Work

Linear Discriminant Analysis (LDA) and Principal Component Analysis (PCA) are classic examples of using linear transformations for projecting data into more interpretable low dimensional spaces. Unsupervised PCA seeks to identify a set of axes that best explain the variance contained in the data. LDA takes a supervised approach, minimizing the intra-class variance and maximizing the inter-class variance given class labeled data points.

Much of the recent work in Distance Metric Learning has focused on learning Mahalanobis distances on the basis of pairwise similarity/dissimilarity constraints. These methods have the same goals as LDA; pairs of points labeled "similar" should be close to one another while pairs labeled "dissimilar" should be distant. MMC [23], a method for identifying a Mahalanobis metric for clustering with side information, uses semidefinite programming to identify a metric that maximizes the sum of distances between points labeled with different classes subject to the constraint that the sum of distances between all points with similar labels be less than some constant.

Large Margin Nearest Neighbor (LMNN) [22] similarly uses semidefinite programming to identify a Mahalanobis distance. In this setting, the algorithm minimizes the sum of distances between a given point and its similarly labeled neighbors while forcing differently labeled neighbors outside of its neighborhood. This method has been shown to be computationally efficient [21] and, in contrast to the similarly motivated Neighborhood Component Analysis [10], is guaranteed to converge to a globally optimal solution. Information Theoretic Metric Learning (ITML) [7] is another popular Distance Metric Learning technique. ITML minimizes the Kullback-Liebler divergence between an initial guess of the matrix that parameterizes the Mahalanobis distance and a solution that satisfies a set of constraints. For surveys of the vast metric learning literature, see [15, 1, 24].

In a dynamic environment, it is necessary to track the changing metric at different times, computing a sequence of estimates of the metric, and to be able to compute those estimates online. Online learning [5] meets these criteria by efficiently updating the estimate every time a new data point is obtained, instead of solving an objective function formed from the entire dataset. Many online learning methods have regret guarantees, that is, the loss in performance relative to a batch method is provably small [5, 9]. In practice, however, the performance of an online learning method is strongly influenced by the learning rate, which may need to vary over time in a dynamic environment [6, 19, 8].

Adaptive online learning methods attempt to address the learning rate problem by continuously updating the learning rate as new observations become available. For learning static parameters, AdaGrad-style methods [19, 8] perform gradient descent steps with the step size adapted based on the magnitude of recent gradients. Follow the regularized leader (FTRL) type algorithms adapt the regularization to the observations [18]. Recently, a method called Strongly Adaptive Online Learning (SAOL) has been proposed for learning parameters undergoing K discrete changes. SAOL maintains several learners with different learning rates and selects the best one based on recent performance [6]. Several of these adaptive methods have provable regret bounds [18, 14, 13]. These typically guarantee low total regret (i.e. regret from time 0 to time T) at every time [18]. SAOL, on the other hand, attempts to have low *static* regret on every subinterval, as well as low regret overall [6]. This allows tracking of discrete changes, but not slow drift.

The remainder of this paper is structured as follows. In Section 2 we formalize the distance metric tracking problem, and section 3 presents the basic COMID online learner. Section 4 presents our

COMID-SADL algorithm, a method of adaptively combining COMID learners with different learning rates. Strongly adaptive bounds on the dynamic regret are presented in Section 5, and results on both synthetic data and a text review dataset are presented in Section 6. Section 7 concludes the paper.

2 Problem Formulation

Metric learning seeks to learn a metric that encourages data points marked as similar to be close and data points marked as different to be far apart. The time-varying Mahalanobis distance at time t is parameterized by $\mathbf{M}_t$ as

$$d^2_{M_t}(\mathbf{x}, \mathbf{z}) = (\mathbf{x} - \mathbf{z})^T \mathbf{M}_t (\mathbf{x} - \mathbf{z}) \tag{1}$$

where $\mathbf{M}_t \in \mathbb{R}^{n \times n} \succeq 0$.

Suppose a temporal sequence of similarity constraints are given, where each constraint is the triplet $(\mathbf{x}_t, \mathbf{z}_t, y_t)$, $\mathbf{x}_t$ and $\mathbf{z}_t$ are data points in $\mathbb{R}^n$, and the label $y_t = +1$ if the points $\mathbf{x}_t, \mathbf{z}_t$ are similar at time t and $y_t = -1$ if they are dissimilar.

Following [16], we introduce the following margin based constraints:

$$t|y_t = 1 : \; d^2_{M_t}(\mathbf{x}_t, \mathbf{z}_t) \leq \mu - 1; \qquad t|y_t = -1 : \; d^2_{M_t}(\mathbf{x}_t, \mathbf{z}_t) \geq \mu + 1, \tag{2}$$

where μ is a threshold that controls the margin between similar and dissimilar points. A diagram illustrating these constraints and their effect is shown in Figure 1. These constraints are softened by penalizing violation of the constraints with a convex loss function ℓ. This gives a loss function

$$\mathcal{L}(\{\mathbf{M}_t, \mu\}) = \frac{1}{T} \sum_{t=1}^{T} \ell(y_t(\mu - \mathbf{u}_t^T \mathbf{M}_t \mathbf{u}_t)) + \rho r(\mathbf{M}_t) = \frac{1}{T} \sum_{t=1}^{T} f_t(\mathbf{M}_t, \mu), \tag{3}$$

where $\mathbf{u}_t = \mathbf{x}_t - \mathbf{z}_t$, r is the regularizer and ρ the regularization parameter. Kunapuli and Shavlik [16] propose using nuclear norm regularization ($r(\mathbf{M}) = \|\mathbf{M}\|_*$) to encourage projection of the data onto a low dimensional subspace (feature selection/dimensionality reduction), and we have also had success with the elementwise L1 norm ($r(\mathbf{M}) = \|\mathrm{vec}(\mathbf{M})\|_1$). In what follows, we develop an adaptive online method to minimize the loss subject to nonstationary smoothness constraints on the sequence of metric estimates $\mathbf{M}_t$.

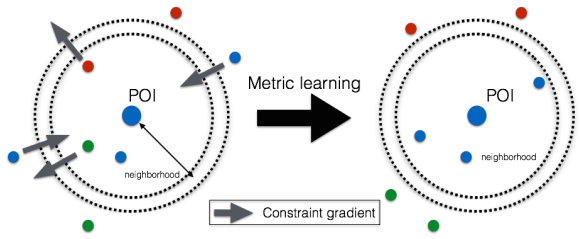

Figure 1: Visualization of the margin based constraints (2), with colors indicating class. The goal of the metric learning constraints is to move target neighbors towards the point of interest (POI), while moving points from other classes away from the target neighborhood.

3 Composite Objective Mirror Descent Update

Viewing the acquisition of new data points as stochastic realizations of the underlying distribution [16] suggests the use of composite objective stochastic mirror descent techniques (COMID).

For the loss (3) and learning rate η_t, COMID [9] gives

$$\begin{aligned} \hat{\mathbf{M}}_{t+1} &= \arg\min_{\mathbf{M} \succeq 0} B_\psi(\mathbf{M}, \hat{\mathbf{M}}_t) + \eta_t \langle \nabla_M \ell_t(\hat{\mathbf{M}}_t, \mu_t), \mathbf{M} - \hat{\mathbf{M}}_t \rangle + \eta_t \rho \|\mathbf{M}\|_* \\ \hat{\mu}_{t+1} &= \arg\min_{\mu \geq 1} B_\psi(\mu, \hat{\mu}_t) + \eta_t \nabla_\mu \ell_t(\hat{\mathbf{M}}_t, \hat{\mu}_t)'(\mu - \hat{\mu}_t), \end{aligned} \tag{4}$$

where B_ψ is any Bregman divergence. In [16] a closed-form algorithm for solving the minimization in (16) with $r(\mathbf{M}) = \|\mathbf{M}\|_*$ is developed for a variety of common losses and Bregman divergences, involving rank one updates and eigenvalue shrinkage.

The output of COMID depends strongly on the choice of η_t. Critically, the optimal learning rate η_t depends on the rate of change of $\mathbf{M}_t$ [11], and thus will need to change with time to adapt to nonstationary drift. Choosing an optimal sequence for η_t is clearly not practical in an online setting with nonstationary drift. We thus introduce COMID-Strongly Adaptive Dynamic Learning (COMID-SADL) as a method to adaptively choose an appropriate learning rate η_t.

4 COMID-SADL

Define a set $\mathcal{I}$ of intervals $I = [t_{I1}, t_{I2}]$ such that the lengths $|I|$ of the intervals are proportional to powers of two, i.e. $|I| = I_0 2^j$, $j = 0, \ldots,$ with an arrangement that is a dyadic partition of the temporal axis. The first interval of length $|I|$ starts at $t = |I|$ (see Figure 2), and additional intervals of length $|I|$ exist such that the rest of time is covered.

Every interval I is associated with a base COMID learner that operates on that interval. Each learner (16) has a constant learning rate proportional to the inverse square of the length of the interval, i.e. $\eta_t(I) = \eta_0/\sqrt{|I|}$. Each learner (besides the coarsest) at level j ($|I| = I_0 2^j$) is initialized to the last estimate of the next coarsest learner (level $j-1$) (see Figure 2). This strategy is equivalent to "backdating" the interval learners so as to ensure appropriate convergence has occurred before the interval of interest is reached, and is effectively a quantized square root decay of the learning rate.

Thus, at a given time t, a set $\text{ACTIVE}(t) \subseteq \mathcal{I}$ of $\text{floor}(\log_2 t)$ intervals/COMID learners are active, running in parallel. Because the metric being learned is changing with time, learners designed for low regret at different scales will have different performance (analogous to the classical bias/variance tradeoff). In other words, there is a scale I_{opt} optimal at a given time.

To select the appropriate scale, we compute weights $w_t(I)$ that are updated based on the learner's recent estimated regret. Our loss function in (3) is unbounded, however, it is a relaxation of an underlying 0-1 loss. For purposes of updating the weights, we propose using a nonlinearity to create a 0-1 loss with a smooth transitions scaled by a parameter c. We choose a linear transition as the nonlinearity

$$\ell_{t,c}(x_t|M_t, \mu_t) = \frac{1}{c}\min\{c, \ell_t(x_t|M_t, \mu_t)\}. \tag{5}$$

We found that using a logistic nonlinearity also gave good results. We set $c = 2$ in all our experiments. The weight update, inspired by the multiplicative weight (MW) literature, is given by

$$\begin{aligned} r_t(I) =& \left(\sum_I \frac{w_t(I)}{\sum_I w_t(I)} \ell_{t,c}(\mathbf{M}_t(I), \mu_t(I))\right) - \ell_{t,c}(\mathbf{M}_t(I), \mu_t(I)) \\ w_{t+1}(I) =& w_t(I)(1 + \eta_I r_t(I)), \quad \forall t \in I. \end{aligned} \tag{6}$$

These hold for all $I \in \mathcal{I}$, where $\eta_I = \min\{1/2, 1/\sqrt{|I|}\}$, $\mathbf{M}_t(I), \mu_t(I)$ are the outputs at time t of the learner on interval I, and $r_t(I)$ is called the estimated regret of the learner on interval I at time t. The initial value of $w(I)$ is η_I. Essentially, this is highly weighting low loss learners and lowly weighting high loss learners.

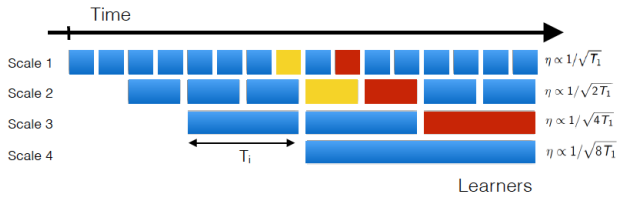

Figure 2: COMID-SADL - Learners at multiple scales run in parallel. Recent observed losses for each learner are used to create weights used to select the appropriate scale at each time. Each yellow and red learner is initialized by the output of the previous learner of the same color, that is, the learner of the next shorter scale.

For any given time t, the output of the learner of interval $I \in \text{ACTIVE}(t)$ is randomly selected as the output of COMID-SADL with probability

$$\Pr(\hat{I}_t = I) = \frac{w_t(I)}{\sum_{I \in \text{ACTIVE}(t)} w_t(I)}. \tag{7}$$

COMID-SADL is summarized in Algorithm 1.

Algorithm 1 COMID-SADL

Initialize: $w_1(I)$
for $t = 1$ to T **do**
  Initialize new learner if needed. New learner at scale $j > 0$: initialize to the last estimate of learner at scale $j - 1$.
  Choose $\hat{I} \in \text{ACTIVE}(t)$ according to (7).
  COMID update (16) for all active learners.
  Set $\mathbf{M}_t \leftarrow \mathbf{M}_t(\hat{I})$, $\mu_t \leftarrow \mu_t(\hat{I})$
  Obtain constraint $(\mathbf{x}_t, \mathbf{z}_t, y_t)$, compute loss $\ell_{t,c}(\cdot)$.
  for $I \in \text{ACTIVE}(t)$ **do**
    Compute estimated regret $r_t(I)$ (6) and update weights: $w_{t+1}(I) = w_t(I)(1 + \eta_I r_t(I))$.
  end for
end for
Return $\{\mathbf{M}_t, \mu_t\}$.

5 Strongly Adaptive Dynamic Regret

The standard static regret is defined as

$$R_{\mathcal{B},static}(I) = \sum_{t \in I} f_t(\hat{\theta}_t) - \min_{\theta \in \Theta} \sum_{t \in I} f_t(\theta). \tag{8}$$

where $f_t(\theta_t)$ is a loss with parameter θ_t. Since in our case the optimal parameter value θ_t is changing, the static regret of an algorithm $\mathcal{B}$ on an interval I is not useful. Instead, let $\mathbf{w} = \{\theta_t\}_{t \in [0,T]}$ be an arbitrary sequence of parameters. Then, the *dynamic regret* of an algorithm $\mathcal{B}$ relative to a comparator sequence $\mathbf{w}$ on the interval I is defined as

$$R_{\mathcal{B},\mathbf{w}}(I) = \sum_{t \in I} f_t(\hat{\theta}_t) - \sum_{t \in I} f_t(\theta_t), \tag{9}$$

where $\hat{\theta}_t$ are generated by $\mathcal{B}$. This allows for a dynamically changing estimate.

In [11] the authors derive dynamic regret bounds that hold over all possible sequences $\mathbf{w}$ such that $\sum_{t \in I} \|\theta_{t+1} - \theta_t\| \leq \gamma$, i.e. bounding the total amount of variation in the estimated parameter. Without this temporal regularization, minimizing the loss would cause θ_t to grossly overfit. In this sense, setting the comparator sequence $\mathbf{w}$ to the "ground truth sequence" or "batch optimal sequence" both provide meaningful intuitive bounds.

Strongly adaptive regret bounds [6] have claimed that static regret is low on every subinterval, instead of only low in the aggregate. We use the notion of dynamic regret to introduce strongly adaptive dynamic regret bounds, proving that *dynamic regret is low on every subinterval $I \subseteq [0, T]$ simultaneously*. In the supplementary material, we prove the following:

Theorem 1 (Strongly Adaptive Dynamic Regret)**.** *Let $\mathbf{w} = \{\mathbf{M}_t\}_{t \in [0,T]}$ be any arbitrary sequence of metrics on the interval $[0, T]$, and define $\gamma_{\mathbf{w}}(I) = \sum_{t \in I} \|\mathbf{M}_{t+1} - \mathbf{M}_t\|$. Then COMID-SADL (Algorithm 1) satisfies*

$$E[R_{COMID-SADL,\mathbf{w}}(I)] \leq \frac{4}{2^{1/2} - 1} C(1 + \gamma_{\mathbf{w}}(I))\sqrt{|I|} + 40 \log(s + 1)\sqrt{|I|}, \tag{10}$$

for every subinterval $I = [q, s] \subseteq [0, T]$ simultaneously. C is a constant, and the expectation is with respect to the random output of the algorithm.

In a dynamic setting, bounds of this type are particularly desirable because they allow for changing *drift rate* and guarantee quick recovery from *discrete changes*. For instance, suppose K discrete switches (large parameter changes or changes in drift rate) occur at times t_i satisfying $0 = t_0 < t_1 < \cdots < t_K = T$. Then since $\sum_{i=1}^{K} \sqrt{|t_{i-1} - t_i|} \leq \sqrt{KT}$, this implies that the total expected dynamic regret on $[0, T]$ remains low ($O(\sqrt{KT})$), while simultaneously guaranteeing that an appropriate learning rate is used on each subinterval $[t_i, t_{i+1}]$.

6 Results

6.1 Synthetic Data

We run our metric learning algorithms on a synthetic dataset undergoing different types of simulated metric drift. We create a synthetic 2000 point dataset with 2 independent 50-20-30% clusterings (A and B) in disjoint 3-dimensional subspaces of $\mathbb{R}^{25}$. The clusterings are formed as 3-D Gaussian blobs, and the remaining 19-dimensional subspace is filled with iid Gaussian noise.

We create a scenario exhibiting nonstationary drift, combining continuous drifts and shifts between the two clusterings (A and B). To simulate continuous drift, at each time step we perform a small random rotation of the dataset. The drift profile is shown in 3. For the first interval, partition A is used and the dataset is static, no drift occurs. Then, the partition is changed to B, followed by an interval of first moderate, then fast, and then moderate drift. Finally, the partition reverts back to A, followed by slow drift.

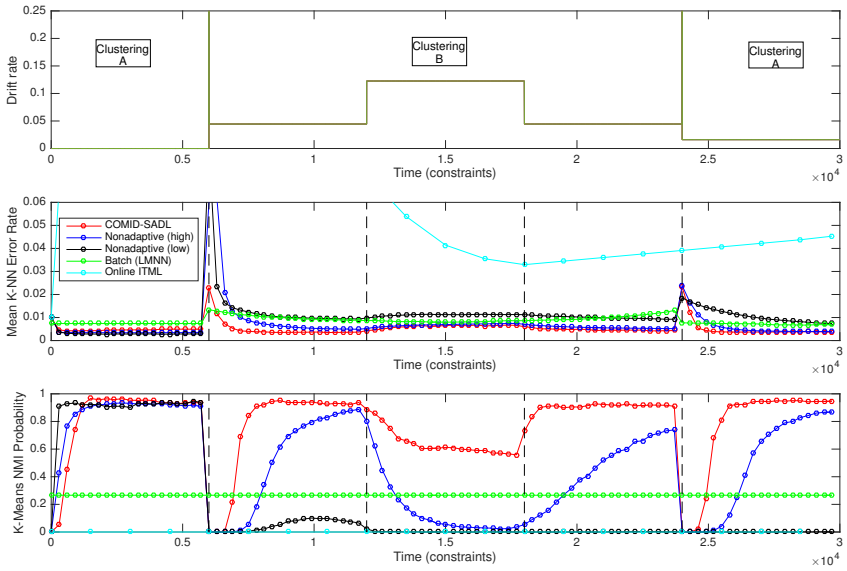

Figure 3: Tracking of a changing metric. Top: Rate of change (scaled Frobenius norm per tick) of the generating metric as a function of time. The large changes result from a change in clustering labels. Metric tracking performance is computed for COMID-SADL (adaptive), nonadaptive COMID (high learning rate), nonadaptive COMID (low learning rate), the batch solution (LMNN), and online ITML, and averaged over 3000 random trials. Shown as a function of time is the mean k-NN error rate (middle) and the probability that the k-means NMI exceeds 0.8 (bottom). Note that COMID-SADL alone is able to effectively adapt to the variety of discrete changes and changes in drift rate.

We generate a series of T constraints from random pairs of points in the dataset, incorporating the simulated drift, running each experiment with 3000 random trials. For each experiment conducted in this section, we evaluate performance using two metrics. We plot the K-nearest neighbor error rate, using the learned embedding at each time point, averaging over all trials. We quantify the clustering performance by plotting the empirical probability that the normalized mutual information (NMI)

of the K-means clustering of the unlabeled data points in the learned embedding at each time point exceeds 0.8 (out of a possible 1). We believe clustering NMI, rather than k-NN performance, is a more realistic indicator of metric learning performance, at least in the case where finding a relevant embedding is the primary goal.

In our results, we consider both COMID-SADL, nonadaptive COMID [16], LMNN (batch) [22], and online ITML [7]. All parameters were set via cross validation. For nonadaptive COMID, we set the high learning rate using cross validation for moderate drift, and we set the low learning rate via cross validation in the case of no drift. The results are shown in Figure 3. Online ITML fails due to its bias agains low-rank solutions [7], and the batch method and low learning rate COMID fail due to an inability to adapt. The high learning rate COMID does well at first, but as it is optimized for slow drift it cannot adapt to the changes in drift rate as well or recover quickly from the two partition changes. COMID-SADL, on the other hand, adapts well throughout the entire interval as expected.

6.2 Clustering Product Reviews

As an example real data task, we consider clustering Amazon text reviews, using the Multi-Domain Sentiment Dataset [4]. We use the 11402 reviews from the Electronics and Books categories, and preprocess the data by computing word counts for each review and 2369 commonly occurring words, thus creating 11402 data points in $\mathbb{R}^{2369}$. Two possible clusterings of the reviews are considered: product category (books or electronics) and sentiment (positive: star rating 4/5 or greater, or negative: 2/5 or less).

Figures 4 and 5 show the first two dimensions of the embeddings learned by static COMID for the category and sentiment clusterings respectively. Also shown are the 2-dimensional standard PCA embeddings, and the k-NN classification performance both before embedding and in each embeddings. As expected, metric learning is able to find embeddings with improved class separability. We emphasize that while improvements in k-NN classification are observed, we use k-NN merely as a way to quantify the separability of the classes in the learned embeddings. In these experiments, we set the regularizer $r(\cdot)$ to the elementwise L1 norm to encourage sparse features.

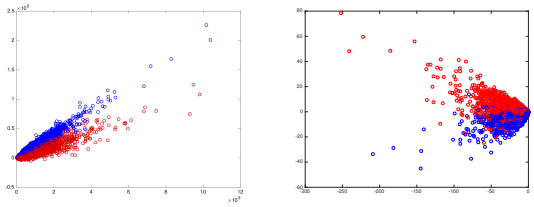

Figure 4: Metric learning for product type clustering. Book reviews blue, electronics reviews red. Original LOO k-NN error rate 15.3%. Left: First two dimensions of learned COMID-SADL embedding (LOO k-NN error rate 11.3%). Right: embedding from PCA (k-NN error 20.4%). Note improved separation of the clusters using COMID-SADL.

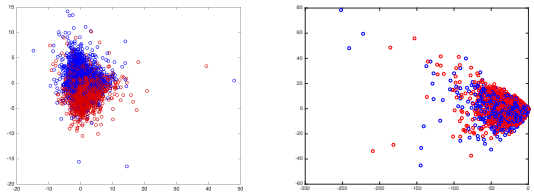

Figure 5: Metric learning for sentiment clustering. Positive reviews blue, negative red. Original LOO k-NN error rate 35.7%. Left: First two dimensions of learned COMID-SADL embedding (LOO k-NN error rate 23.5%). Right: embedding from PCA (k-NN error 41.9%). Note improved separation of the clusters using COMID-SADL.

We then conducted drift experiments where the clustering changes. The change happens after the metric learner for the original clustering has converged, hence the nonadaptive learning rate is

effectively zero. For each change, we show the k-NN error rate in the learned COMID-SADL embedding as it adapts to the new clustering. Emphasizing the visualization and computational advantages of a low-dimensional embedding, we computed the k-NN error after projecting the data into the first 5 dimensions of the embedding. Also shown are the results for a learner where an oracle allows reinitialization of the metric to the identity at time zero, and the nonadaptive learner for which the learning rate is not increased. Figure 6 (left) shows the results when the clustering changes from the four class sentiment + type partition to the two class product type only partition, and Figure 6 (right) shows the results when the partition changes from sentiment to product type. In the first case, the similar clustering allows COMID-SADL to significantly outperform even the reinitialized method, and in the second remain competitive where the clusterings are unrelated.

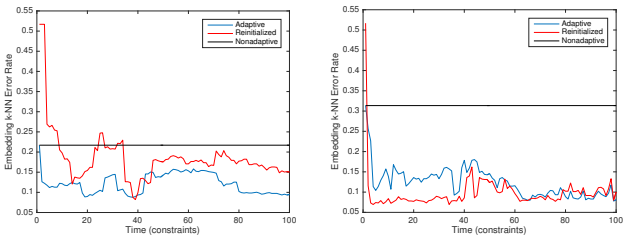

Figure 6: Metric drift in Amazon review data. Left: Change from product type + sentiment clustering to simply product type; Right: Change from sentiment to product type clustering. The proposed COMID-SADL adapts to changes, tracking the clusters as they evolve. The oracle reinitialized mirror descent method (COMID) learner has higher tracking error and the nonadaptive learner (straight line) does not track the changes at all.

7 Conclusion and Future Work

Learning a metric on a complex dataset enables both unsupervised methods and/or a user to home in on the problem of interest while de-emphasizing extraneous information. When the problem of interest or the data distribution is nonstationary, however, the optimal metric can be time-varying. We considered the problem of tracking a nonstationary metric and presented an efficient, strongly adaptive online algorithm, called COMID-SADL, that has strong theoretical regret guarantees. Performance of our algorithm was evaluated both on synthetic and real datasets, demonstrating its ability to learn and adapt quickly in the presence of changes both in the clustering of interest and in the underlying data distribution.

Potential directions for future work include the learning of more expressive metrics beyond the Mahalanobis metric, the incorporation of unlabeled data points in a semi-supervised learning framework [2], and the incorporation of an active learning framework to select which pairs of data points to obtain labels for at any given time [20].

8 Acknowledgments

The Lincoln Laboratory portion of this work was sponsored by the Assistant Secretary of Defense for Research and Engineering under Air Force Contract #FA8721-05-C-0002. Opinions, interpretations, conclusions and recommendations are those of the author and are not necessarily endorsed by the United States Government.

9 Strongly Adaptive Dynamic Regret

We will prove the following general theorem giving strongly adaptive dynamic regret bounds.

Theorem 2. *Let* $\mathbf{w} = \{\theta_1, \ldots, \theta_T\}$ *be an arbitrary sequence of parameters and define* $\gamma_{\mathbf{w}}(I) = \sum_{q \le t < s} \|\theta_{t+1} - \theta_t\|$ *as a function of* $\mathbf{w}$ *and an interval* $I = [q, s]$*. Choose a set of learners* $\mathcal{B}$ *such that given an interval* I *the learner* $\mathcal{B}_I$ *satisfies*

$$R_{\mathcal{B}_I, \mathbf{w}}(T) \le C(1 + \gamma_{\mathbf{w}}(I))\sqrt{|I|} \tag{11}$$

for some constant $C > 0$. Then the strongly adaptive dynamic learner $SADL^{\mathcal{B}}$ (COMID-SADL) using $\mathcal{B}$ as the interval learners satisfies

$$R_{SADL^{\mathcal{B}},\mathbf{w}}(I) \leq 8C(1+\gamma(I))\sqrt{|I|} + 40\log(s+1)\sqrt{|I|} \tag{12}$$

on every interval $I = [q, s] \subseteq [0, T]$.

The proof techniques are similar to those found in [6**?**] which are in turn similar to the analysis of the Multiplicative Weights Update (MW) method.

Define

$$\begin{aligned} \tilde{w}_t &= \begin{cases} 0 & t < q \\ 1 & t = q \\ \tilde{w}_{t-1}(I)(1+\eta_I r_{t-1}(I)) & q < t \leq s+1 \\ \tilde{w}_s(I) & t > s+1 \end{cases} \\ \tilde{W}_t &= \sum_{I \in \mathcal{I}} \tilde{w}_{t+1}(I). \end{aligned} \tag{13}$$

Note that $w_t(I) = \eta_I I(t)\tilde{w}_t(I)$ where $I(t)$ is the indicator function for I.

We first prove a pair of lemmas.

Lemma 1.

$$E[\tilde{W}_t] \leq t(\log(t)+1)$$

for all $t \geq 1$.

Proof. For all $t \geq 1$, $|\{[q,s] \in \mathcal{I} : q = t\}| \leq \lfloor \log(t) \rfloor + 1$. Thus

$$\begin{aligned} \tilde{W}_{t+1} &= \sum_{I=[q,s]\in\mathcal{I}} \tilde{w}_{t+1}(I) \\ &= \sum_{I=[t+1,s]\in\mathcal{I}} \tilde{w}_{t+1}(I) + \sum_{I=[q,s]\in\mathcal{I}:q\leq t} \tilde{w}_{t+1}(I) \\ &\leq \log(t+1) + 1 + \sum_{I=[q,s]\in\mathcal{I}:q\leq t} \tilde{w}_{t+1}(I). \end{aligned}$$

Then

$$\begin{aligned} \sum_{I=[q,s]\in\mathcal{I}:q\leq t} \tilde{w}_{t+1}(I) &= \sum_{I=[q,s]\in\mathcal{I}:q\leq t} \tilde{w}_t(I)(1+\eta_I I(t) r_t(I)) \\ &= \tilde{W}_t + \sum_{I\in\mathcal{I}} w_t(I) r_t(I). \end{aligned}$$

Suppose that $E[\tilde{W}_t] \leq t(\log(t)+1)$. Furthermore, note that

$$\begin{aligned} E\left[\sum_{I\in\mathcal{I}} w_t(I) r_t(I)\right] &= W_t \sum_{I\in\mathcal{I}} p_t(I)(E[\ell_t(x_t)] - \ell_t(x_t(I))) \\ &= E[\ell_t(x_t)] - \sum_{I\in\mathcal{I}} p_t(I)\ell_t(x_t(I))) \\ &= E[\ell_t(x_t)] - E[\ell_t(x_t)] \\ &= 0. \end{aligned}$$

since $x_t = x_t(I)$ with probability $p_t(I)$. Thus

$$\begin{aligned} E[\tilde{W}_{t+1}] &\leq t(\log(t)+1) + \log(t+1) + 1 + E\left[\sum_{I\in\mathcal{I}} w_t(I) r_t(I)\right] \\ &\leq (t+1)(\log(t+1)+1). \end{aligned}$$

Since $E[\tilde{W}_1] = \tilde{W}_1 = \tilde{w}([1,1]) = 1$, the lemma follows by induction.

□

Lemma 2.

$$E\sum_{t=q}^{s} r_t(I) \leq 5\log(s+1)\sqrt{|I|},$$

for every $I = [q, s] \in \mathcal{I}$.

Proof. Fix $I = [q, s] \in \mathcal{I}$. Recall that

$$\tilde{w}_{s+1}(I) = \prod_{t=q}^{s}(1 + \eta_I I(t) r_t(I)) = \prod_{t=q}^{s}(1 + \eta_I r_t(I)).$$

Since $\eta_I \in (0, 1/2)$ and $\log(1 + x) \geq (x - x^2)$ for all $x \geq -1/2$,

$$\begin{aligned}\log(\tilde{w}_{s+1}(I)) &= \sum_{t=q}^{s}\log(1 + \eta_I r_t(I)) \\ &\geq \sum_{t=q}^{s}\eta_I r_t(I) - \sum_{t=q}^{s}(\eta_I r_t(I))^2 \\ &\geq \eta_I\left(\sum_{t=q}^{s} r_t(I) - \eta_I|I|\right).\end{aligned} \tag{14}$$

By Lemma 1 we have

$$E[\tilde{w}_{s+1}(I)] \leq E[\tilde{W}_{s+1}] \leq (s+1)(\log(s+1)+1),$$

so

$$E[\log(\tilde{w}_{s+1}(I))] \leq \log(E[\tilde{w}_{s+1}]) \leq \log(s+1) + \log(\log(s+1)+1).$$

Combining with the expectation of (14) and dividing by η_I,

$$\begin{aligned}E\left[\sum_{t=q}^{s} r_t(I)\right] &\leq \eta_I|I| + \frac{1}{\eta_I}(\log(s+1) + \log(\log(s+1)+1)) \\ &\leq \eta_I|I| + 2\eta_I^{-1}\log(s+1) \\ &= 5\log(s+1)\sqrt{|I|}.\end{aligned}$$

since $x \geq \log(1+x)$ and $\eta_I = \min\{1/2, |I|^{-1/2}\}$.

□

Define the restriction of $\mathcal{I}$ to an interval $J \subseteq \mathbb{N}$ as $\mathcal{I}|_J = \{I \in \mathcal{I} : I \subseteq J\}$. Note the following lemma from [6]:

Lemma 3. *Consider the arbitrary interval $I = [q, s] \subseteq \mathbb{N}$. Then, the interval I can be partitioned into two finite sequences of disjoint and consecutive intervals, given by $(I_{-k}, \ldots, I_0) \subseteq \mathcal{I}|_I$ and $(I_1, I_2, \ldots, I_p) \subseteq \mathcal{I}|_I$, such that*

$$\begin{aligned}&|I_{-i}|/|I_{-i+1}| \leq 1/2, \quad \forall i \geq 1, \\ &|I_i|/|I_{i-1}| \leq 1/2, \quad \forall i \geq 2.\end{aligned}$$

This enables us to extend the bounds to every arbitrary interval $I = [q, s] \subseteq [T]$ and thus complete the proof.

Let $I = \bigcup_{i=-k}^{p} I_i$ be the partition described in Lemma 3. Then

$$R_{SADL^{\mathcal{B}},\mathbf{w}}(I) \leq \sum_{i\leq 0} R_{SADL^{\mathcal{B}},\mathbf{w}}(I_i) + \sum_{i\geq 1} R_{SADL^{\mathcal{B}},\mathbf{w}}(I_i). \tag{15}$$

By Lemma 2 and (11),

$$\sum_{i\leq 0} R_{SADL^{\mathcal{B}},\mathbf{w}}(I_i) \leq C\sum_{i\leq 0}(1+\gamma_{\mathbf{w}}(I_i))\sqrt{|I_i|} + 5\sum_{i\leq 0}\log(s_i+1)\sqrt{I_i}$$
$$\leq (C(1+\gamma(I)) + 5\log(s_i+1))\sum_{i\leq 0}\sqrt{I_i},$$

since $\gamma_{\mathbf{w}}(I_i) \leq \gamma_{\mathbf{w}}(I)$ by definition. By Lemma 3,

$$\sum_{i\leq 0}\sqrt{|I_i|} \leq \frac{\sqrt{2}}{\sqrt{2}-1}\sqrt{|I|} \leq 4\sqrt{|I|}.$$

This bounds the first term of the right hand side of Equation (15). The bound for the second term can be found in the same way. Thus,

$$R_{SADL^{\mathcal{B}},\mathbf{w}}(I) \leq (8C(1+\gamma(I))\sqrt{|I|} + 40\log(s+1)\sqrt{|I|}.$$

Since this holds for all I, this completes the proof.

10 Online DML Dynamic Regret

In this section, we derive the dynamic regret of the COMID metric learning algorithm. Recall that the COMID algorithm is given by

$$\begin{aligned}\hat{\mathbf{M}}_{t+1} &= \arg\min_{\mathbf{M}\succeq 0} B_\psi(\mathbf{M}, \hat{\mathbf{M}}_t) \\ &\quad + \eta_t\langle\nabla_M \ell_t(\hat{\mathbf{M}}_t, \mu_t), \mathbf{M} - \hat{\mathbf{M}}_t\rangle + \eta_t\rho\|\mathbf{M}\|_* \\ \hat{\mu}_{t+1} &= \arg\min_{\mu\geq 1} B_\psi(\mu, \hat{\mu}_t) + \eta_t\nabla_\mu \ell_t(\hat{\mathbf{M}}_t, \hat{\mu}_t)'(\mu - \hat{\mu}_t),\end{aligned} \tag{16}$$

where B_ψ is any Bregman divergence and η_t is the learning rate parameter. From [11] we have:

Theorem 3.

$$G_\ell = \max_{\theta\in\Theta,\ell\in\mathcal{L}}\|\nabla f(\theta)\|$$
$$\phi_{max} = \frac{1}{2}\max_{\theta\in\Theta}\|\nabla\psi(\theta)\|$$
$$D_{max} = \max_{\theta,\theta'\in\Theta} B_\psi(\theta'\|\theta)$$

Let the sequence $\hat{\theta}_t = [\hat{\mathbf{M}}_t, \hat{\mu}_t]$, $t = 1, \cdots, T$ be generated via the COMID algorithm, and let $\mathbf{w}$ be an arbitrary sequence in $\mathcal{W} = \{\mathbf{w}|\sum_{t=0}^{T-1}\|\theta_{t+1} - \theta_t\| \leq \gamma\}$. Then using $\eta_{t+1} \leq \eta_t$ gives a dynamic regret

$$R_{\mathbf{w}}([0,T]) \leq \frac{D_{max}}{\eta_{T+1}} + \frac{4\phi_{max}}{\eta_T}\gamma + \frac{G_\ell^2}{2\sigma}\sum_{t=1}^{T}\eta_t \tag{17}$$

Using a nonincreasing learning rate η_t, we can then prove a bound on the dynamic regret for a quite general set of stochastic optimization problems.

Applying this to our problem, we have

$$G_\ell = \max_{\|\mathbf{M}\|\leq c,t,\mu}\|\nabla(\ell_t(\mathbf{M},\mu) + \rho\|\mathbf{M}\|_*)\|_2$$
$$\phi_{max} = \frac{1}{2}\max_{\|\mathbf{M}\|\leq c}\|\nabla\psi(\mathbf{M})\|_2$$
$$D_{max} = \max_{\|\mathbf{M}\|,\|\mathbf{M}'\|\leq c} B_\psi(\mathbf{M}'\|\mathbf{M})$$

For $\ell_t(\cdot)$ being the hinge loss and $\psi = \|\cdot\|_F^2$,

$$\begin{aligned} G_\ell &\le \sqrt{(\max_t \|\mathbf{x}_t - \mathbf{z}_t\|_2^2 + \rho)^2} \\ \phi_{max} &= c\sqrt{n} \\ D_{max} &= 2c\sqrt{n}. \end{aligned}$$

The other two quantities are guaranteed to exist and depend on the choice of Bregman divergence and c. Thus,

Corollary 1 (Dynamic Regret: ML COMID)**.** *Let the sequence $\hat{\mathbf{M}}_t, \hat{\mu}_t$ be generated by* (16)*, and let $\mathbf{w} = \{\mathbf{M}_t\}_{t=1}^T$ be an arbitrary sequence with $\|\mathbf{M}_t\| \le c$. Then using $\eta_{t+1} \le \eta_t$ gives*

$$R_{\mathbf{w}} \le \frac{D_{max}}{\eta_{T+1}} + \frac{4\phi_{max}}{\eta_T}\gamma + \frac{G_\ell^2}{2\sigma}\sum_{t=1}^T \eta_t \tag{18}$$

and setting $\eta_t = \eta_0/\sqrt{T}$,

$$\begin{aligned} R_{\mathbf{w}}([0,T]) \le& \sqrt{T}\left(\frac{D_{max} + 4\phi_{max}(\sum_t \|\mathbf{M}_{t+1} - \mathbf{M}_t\|_F)}{\eta_0} + \frac{\eta_0 G_\ell^2}{2\sigma}\right) \\ =& O\left(\sqrt{T}\left[1 + \sum_{t=1}^T \|\mathbf{M}_{t+1} - \mathbf{M}_t\|_F\right]\right). \end{aligned} \tag{19}$$

Corollary 1 is a bound on the regret relative to the batch estimate of $\mathbf{M}_t$ that minimizes the total batch loss subject to a bounded variation $\sum_t \|\mathbf{M}_{t+1} - \mathbf{M}_t\|_F$. Also note that setting $\eta_t = \eta_0/\sqrt{t}$ gives the same bound as (19).

In other words, we pay a linear penalty on the total amount of variation in the underlying parameter sequence. From (19), it can be seen that the bound-minimizing η_0 increases with increasing $\sum_t \|\mathbf{M}_{t+1} - \mathbf{M}_t\|_F$, indicating the need for an adaptive learning rate.

For comparison, if the metric is in fact static then by standard stochastic mirror descent results [11]

Theorem 4 (Static Regret)**.** *If $\hat{\mathbf{M}}_1 = 0$ and $\eta_t = (2\sigma D_{max})^{1/2}/(G_f\sqrt{T})$, then*

$$R_{static}([0,T]) \le G_f(2TD_{max}/\sigma)^{1/2}. \tag{20}$$

11 COMID-SADL Bound

Let $\mathcal{B}^i$ be a COMID learner at any of the scales used in $SADL^{\mathcal{B}}$, with output $x_t(i)$. Define the relative regret

$$\tilde{R}^i_{SADL,\mathbf{w}}(I) = \sum_{t\in I} \ell_c(x_t) - \ell_c(x_t(i))$$

as the extra ℓ_c loss suffered relative to the algorithm $\mathcal{B}^i$. From the proof of Theorem 1 we have

Lemma 4. *For any c, the following holds simultaneously for all $\mathcal{B}^i$ and I.*

$$\tilde{R}^i_{SADL,\mathbf{w}}(I) \le 40\log(s+1)|I|^{1/2}.$$

This implies that SADL incurs at most $O(\sqrt{|I|})$ additional scaled 0-1 loss on any interval relative to each of the base learners, all of which have low regret in the convex loss. Due to the nonconvexity of the scaled 0-1 loss, it is difficult to state more for arbitrary c.

However, since $\|\mathbf{M}\| \le c'$, $\ell_t(\mathbf{M}_t, \mu_t) \le k = \ell(c' \max_t \|\mathbf{x}_t - \mathbf{z}_t\|_2^2)$. Hence for $c = k$, $\ell_c(\cdot) = \frac{1}{c}\ell(\cdot)$ everywhere. Thus Corollary 1 can be used in Theorem 2, giving

Theorem 5 (COMID-SADL)**.** *Let $\mathcal{B}$ be the COMID algorithm of* (16) *with $\eta_t(I) = \eta_0/\sqrt{|I|}$. Then there exists a c such that the strongly adaptive online learner $SADL^{\mathcal{B}}$ (COMID-SADL) satisfies*

$$R_{SADL,\mathbf{w}}(I) \le 8C(1 + \gamma_{\mathbf{w}}(I))|I|^{1/2} + 40\log(s+1)|I|^{1/2} \tag{21}$$

for some constant C and every interval $I = [q, s]$.

Note that this bound also holds for the original convex loss ℓ.

References

[1] Aurélien Bellet, Amaury Habrard, and Marc Sebban. A survey on metric learning for feature vectors and structured data. *arXiv preprint arXiv:1306.6709*, 2013.

[2] Mikhail Bilenko, Sugato Basu, and Raymond J Mooney. Integrating constraints and metric learning in semi-supervised clustering. In *ICML*, page 11, 2004.

[3] Christopher M Bishop. *Pattern Recognition and Machine Learning*. Springer, 2006.

[4] John Blitzer, Mark Dredze, Fernando Pereira, et al. Biographies, bollywood, boom-boxes and blenders: Domain adaptation for sentiment classification. In *ACL*, volume 7, pages 440–447, 2007.

[5] Nicolo Cesa-Bianchi and Gábor Lugosi. *Prediction, learning, and games*. Cambridge University Press, 2006.

[6] Amit Daniely, Alon Gonen, and Shai Shalev-Shwartz. Strongly adaptive online learning. *ICML*, 2015.

[7] Jason V Davis, Brian Kulis, Prateek Jain, Suvrit Sra, and Inderjit S Dhillon. Information-theoretic metric learning. In *ICML*, pages 209–216, 2007.

[8] John C Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. In *COLT*, 2010.

[9] John C Duchi, Shai Shalev-Shwartz, Yoram Singer, and Ambuj Tewari. Composite objective mirror descent. In *COLT*, pages 14–26. Citeseer, 2010.

[10] Jacob Goldberger, Geoffrey E Hinton, Sam T Roweis, and Ruslan Salakhutdinov. Neighbourhood components analysis. In *Advances in neural information processing systems*, pages 513–520, 2004.

[11] E.C. Hall and R.M. Willett. Online convex optimization in dynamic environments. *Selected Topics in Signal Processing, IEEE Journal of*, 9(4):647–662, June 2015.

[12] Trevor Hastie, Robert Tibshirani, Jerome Friedman, and James Franklin. The elements of statistical learning: data mining, inference and prediction. *The Mathematical Intelligencer*, 27 (2):83–85, 2005.

[13] Elad Hazan and C Seshadhri. Adaptive algorithms for online decision problems. In *Electronic Colloquium on Computational Complexity (ECCC)*, volume 14, 2007.

[14] Mark Herbster and Manfred K Warmuth. Tracking the best expert. *Machine Learning*, 32(2): 151–178, 1998.

[15] Brian Kulis. Metric learning: A survey. *Foundations and Trends in Machine Learning*, 5(4): 287–364, 2012.

[16] Gautam Kunapuli and Jude Shavlik. Mirror descent for metric learning: a unified approach. In *Machine Learning and Knowledge Discovery in Databases*, pages 859–874. Springer, 2012.

[17] John A Lee and Michel Verleysen. *Nonlinear dimensionality reduction*. Springer Science & Business Media, 2007.

[18] H Brendan McMahan. Analysis techniques for adaptive online learning. *arXiv preprint arXiv:1403.3465*, 2014.

[19] H Brendan McMahan and Matthew Streeter. Adaptive bound optimization for online convex optimization. In *COLT*, 2010.

[20] Burr Settles. Active learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 6(1):1–114, 2012.

[21] Kilian Q Weinberger and Lawrence K Saul. Fast solvers and efficient implementations for distance metric learning. In *ICML*, pages 1160–1167, 2008.

[22] Kilian Q Weinberger, John Blitzer, and Lawrence K Saul. Distance metric learning for large margin nearest neighbor classification. In *Advances in Neural Information Processing System*, pages 1473–1480, 2005.

[23] Eric P Xing, Michael I Jordan, Stuart Russell, and Andrew Y Ng. Distance metric learning with application to clustering with side-information. In *Advances in Neural Information Processing Systems*, pages 505–512, 2002.

[24] Liu Yang and Rong Jin. Distance metric learning: A comprehensive survey. *Michigan State Universiy*, 2, 2006.